

TWO-GENERATOR SUBGROUPS OF RIGHT-ANGLED ARTIN GROUPS ARE QUASI-ISOMETRICALLY EMBEDDED

MIKE CARR

ABSTRACT. We show that non-abelian two-generator subgroups of right-angled Artin groups are quasi-isometrically embedded free groups. This provides an alternate proof of a theorem of A. Baudisch: that all two-generator subgroups are free or free abelian. Additionally, it shows that they are quasi-isometrically embedded. Our theorem also gives a method for detecting groups that are not isomorphic to a subgroup of any RAAG. We present some counterexamples in subgroups with more than two generators.

1. INTRODUCTION

Definition 1.1. Given a combinatorial graph Γ with vertex set V and edge set E , the *right-angled Artin group* A_Γ is the group presented by the generating set V and the relations $\{v_i v_j = v_j v_i \mid (v_i, v_j) \in E\}$.

Right-angled Artin groups (RAAGs) constitute a spectrum between free abelian groups, given by complete graphs, and free groups, given by edgeless graphs. The simple definition belies that fact that complicated groups can exist as subgroups of right-angled Artin groups. Crisp and Weist showed that nearly all surface groups can be embedded in some RAAG [CW04]. In addition to providing a wealth of interesting examples, RAAG subgroups have underpinned several recent results in group theory and topology. The recent proof of the virtual Haken conjecture by Ian Agol [AGM12] relied on showing that every hyperbolic 3-manifold group had a finite index subgroup that embeds in a RAAG via a map defined by Haglund and Wise [HW08]. Hsu and Wise also used embeddings of subgroups to show that certain graph groups are linear [HW10].

The behavior of two-generator subgroups of RAAGs is much more circumscribed. A theorem of A. Baudisch completely describes their group structure.

Theorem 1.2. [Bau81, 1.3] *Every two generator subgroup of A_Γ is either free or free abelian.*

This description of two-generator subgroups is interesting for at least two reasons. First, it passes to subgroups. If this description holds for a group G , then it holds for every subgroup $H < G$. Second, two generator subgroups of many other well-studied groups have this description or some approximation of it. In mapping class groups, it is false in general, but given two generators, one can take appropriate powers of each such that the group those powers generate is free or free abelian.

Two generator subgroups of pure braid groups do fit this description. This fact that was used by Scrimshaw [Scr10] and Kim and Koberda [KK13] who recovered Baudisch's result by embedding any A_Γ in some pure braid group.

Mike Carr was partially supported by NSF grant DMS-1106726.

Our main result concerns the metric structure of two-generator subgroups in a right-angled Artin group. We use the word metric to realize each group as a metric space. The desired metric property is the following.

Definition 1.3. A *quasi-isometric embedding* is a map $f : X \rightarrow Y$ between metric spaces such that there are constants $\lambda \geq 1$ and $\epsilon \geq 0$ that satisfy

$$\frac{1}{\lambda}d_X(p, q) - \epsilon \leq d_Y(f(p), f(q)) \leq \lambda d_X(p, q) + \epsilon$$

for all $p, q \in X$.

The metric behavior of abelian subgroups of a RAAG is already well understood, via the Flat Torus Theorem (see in [BH99, II.7.1 and II.7.17]). Our main theorem resolves the nonabelian case.

Theorem 1.4. *Let $F = \langle u, v \rangle$ be a free group. If $\phi : F \rightarrow A_\Gamma$ is a homomorphism, such that $\phi(u)$ and $\phi(v)$ do not commute, then ϕ is a quasi-isometric embedding.*

We retrieve Baudisch’s original, algebraic description (1.2) as a corollary. Combining with the abelian case, we produce the following metric description.

Corollary 1.5. *Every two generator subgroup of A_Γ is a quasi-isometrically embedded free or free abelian group.*

One consequence is that if an injective homomorphism $\psi : F_2 \rightarrow G$ is not quasi-isometric, then G is not isomorphic to a subgroup in any RAAG.

There is also a divergence between the algebraic and metric treatments in higher rank free groups. We produce a free subgroup of a RAAG whose inclusion is not a quasi-isometric embedding.

2. BACKGROUND

This section will recall some standard results from geometric group theory, beginning with a discussion of quasi-isometric embeddings.

2.1. Quasi-isometries. It is a straightforward exercise to show that compositions of quasi-isometric embeddings are quasi-isometric embeddings. The next natural step is to define a relation based on those embeddings that have an inverse, defined as follows.

Definition 2.1. A *quasi-isometry* is a quasi-isometric embedding such that all of Y lies within a fixed distance of the image $f(X)$. If such a map exists, then X and Y are *quasi-isometric*.

One can show that a quasi-isometry f has a quasi-inverse, a quasi-isometry g such that $f \circ g$ and $g \circ f$ move points a bounded distance. “Quasi-isometric” forms an equivalence relation on metric spaces. It is the natural equivalence relation for the study of groups as metric spaces. From any generating set S of a group G , one can produce the word metric d where $d(g_1, g_2)$ is the length of the shortest word in S that represents $g_1^{-1}g_2$. We can’t directly use geodesics to study G , since it has none. Fortunately the Cayley graph of G does have geodesics, and it is straightforward to prove the following:

Proposition 2.2. *If G has generating set S then the Cayley graph $\text{Cay}(G, S)$ (with unit length edges) is quasi-isometric to the group G given the word metric from S .*

The other natural appeal of quasi-isometry is that the word metric depends on the choice of generating set S . Quasi-isometry resolves this ambiguity.

Proposition 2.3. *If two finite generating sets S_1 and S_2 of a group G give word metrics d_1 and d_2 , then the identity map on G is a quasi-isometry between (G, d_1) and (G, d_2) .*

A more general and more powerful theorem concerning quasi-isometries of groups is the following:

Theorem 2.4 (The Švarc-Milnor Lemma [Mil68]). *If a group G acts properly, cocompactly and by isometries on a geodesic metric space X , then G is finitely generated. Furthermore for any $x \in X$, the orbit map $g \mapsto gx$ is a quasi-isometry.*

An accessible proof is found in [BH99, I.8.19].

2.2. CAT(0) spaces. We will denote a geodesic segment between points a and b in a geodesic metric space by $[a, b]$.

Definition 2.5. A geodesic metric space X has the CAT(0) property if for any three points $p, q, r \in X$, any geodesic triangle $[p, q] \cup [q, r] \cup [r, p]$ is at least as thin as a comparison triangle $[P, Q] \cup [Q, R] \cup [R, P]$ with equal length sides in the euclidean plane \mathbb{E} . That is, for any $x \in [p, q]$ and $y \in [q, r]$, the corresponding points $X \in [P, Q]$ and $Y \in [Q, R]$ with $d_X(x, p) = d_{\mathbb{E}}(X, P)$ and $d_X(y, q) = d_{\mathbb{E}}(Y, Q)$ have the property:

$$d_X(x, y) \leq d_{\mathbb{E}}(X, Y).$$

If every point in X has a neighborhood that is CAT(0), then we say X is *non-positively curved*.

The CAT(0) property has many useful consequences. A general exposition is found in [BH99, II.1].

Proposition 2.6. *If X is CAT(0) and $p, q, r \in X$ with $d_X(q, r) < \epsilon$ then the geodesics $[p, q]$ and $[p, r]$ each lie within an ϵ -neighborhood of the other.*

This follows directly from the equivalent fact in the comparison triangle. Here is a powerful application.

Corollary 2.7. *If X is CAT(0) then there is only one geodesic between any two points p and q .*

Corollary 2.8. *A CAT(0) space is contractible.*

The contraction moves each point to the basepoint along the unique geodesic between them at constant speed. Continuity follows directly from Proposition 2.6 and the fact that geodesics are continuous maps.

Proposition 2.9. *In a CAT(0) space, local geodesics are (global) geodesics, and are thus unique.*

Non-positive curvature can also be used to study covering spaces and fundamental groups. One critical ingredient is the following generalization of a theorem from Riemannian geometry, whose name it also shares.

Theorem 2.10 (The Cartan-Hadamard Theorem for Non-Positively Curved Spaces). [AB90, Theorem 1] *If X is non-positively curved, then the universal cover \tilde{X} is $CAT(0)$.*

The following standard application illustrates the connection between metric and group theoretic aspects of maps. A proof is in [BH99, II.4.14].

Proposition 2.11. *If X is non-positively curved, Y is a geodesic metric space and $f : Y \rightarrow X$ is a locally isometric embedding, then the induced map $f_* : \pi_1(Y) \rightarrow \pi_1(X)$ is injective.*

Furthermore, by the Švarc-Milnor Lemma, the map f_* is a quasi-isometric embedding, since it is the composition of two quasi-isometries and an isometric embedding:

$$\pi_1(Y) \rightarrow \tilde{Y} \rightarrow \tilde{X} \rightarrow \pi_1(X)$$

2.3. Isometries of $CAT(0)$ spaces. Like in the hyperbolic plane, isometries of a $CAT(0)$ space are classified as either elliptic, parabolic, or hyperbolic.

If G is a group acting properly by isometries on X such that X/G is compact, then a standard result [BH99, II.6.10.2] is that all $g \in G$ are either elliptic or hyperbolic. If G is a deck group, then it acts freely and all its isometries are hyperbolic.

A hyperbolic isometry of the hyperbolic plane has an axis, a bi-infinite geodesic preserved by the isometry.

If g is a hyperbolic isometry of a $CAT(0)$ space, then it also has an axis. This axis may not be unique, but any two axes lie in bounded neighborhoods of each other. This is clear since g preserves distance. A $CAT(0)$ argument shows that any two axes are actually parallel (see [BH99, II.6.8]).

2.4. Cube complexes.

Definition 2.12. A *cube complex* is a metric space X that is the union of metric cubes $[-\frac{1}{2}, \frac{1}{2}]^n$ glued by isometries of their faces.

Finite dimensional cube complexes are geodesic metric spaces [BH99, I.7.19].

Given a RAAG A_Γ , we take a cube $[-\frac{1}{2}, \frac{1}{2}]^{\#V}$ with the euclidean metric and with opposite faces identified to make an n -torus. The faces of this torus form a complex of metric cubes, specifically, each edge is an interval associated to an element of V . We consider the subcomplex of only those cubes whose edges all commute pairwise in A_Γ .

This subcomplex has fundamental group A_Γ . We call it the *Salvetti complex*, which we will denote X_Γ . With the path metric, X_Γ satisfies Gromov's link condition [Gro87, 4.2c] for a non-positively curved cube complex [CD95, 3.1.1].

In addition, by the Švarc-Milnor Lemma, the universal cover \tilde{X}_Γ is quasi-isometric to A_Γ .

2.5. Hyperplanes.

Definition 2.13. Given a metric n -cube: $[-\frac{1}{2}, \frac{1}{2}]^n$, we define the i^{th} *midcube* to be a set of points that is 0 in the i^{th} coordinate. A midcube naturally inherits the metric and cell structure of an $(n - 1)$ -cube.

In a cube complex X we define a relation on all midcubes in X . Two midcubes are related if and only if they share a vertex. There is a unique weakest equivalence relation generated by this relation. This equates any midcubes that can be connected by a succession of vertices. The union of cubes in an equivalence class is a *hyperplane*. Figure 1 shows examples in a small complex.

We'll say that the vertices of X that belong to edges of X crossing a hyperplane h are *adjacent to h* .




FIGURE 1. 3 hyperplanes in a cube complex. p is adjacent to h_1 and h_2 but not h_3 .

If X is CAT(0) then a hyperplane h contains at most one midcube of a given cube, and divides X into two connected components [Sag95, 4.10]. The cubes containing h form an $h \times I$ neighborhood with a metric $\sqrt{d_h^2 + d_I^2}$. This makes geodesics near h very easy to compute, and since geodesic segments are unique in X we obtain the following standard results:

- (1) If $p, q \in h$, then the geodesic segment from p to q (in X) is also contained in h . Specifically, h is simply connected.
- (2) A geodesic in X is either contained in h , disjoint from h or intersects h transversely at a single point.

We'll make use of the following lemma in our proof of the main theorem.

Lemma 2.14. *If L is a bi-infinite geodesic in a CAT(0) cube complex that transversely intersects a hyperplane h , then the distances from points on L to h is unbounded, traveling in either direction along L .*

Proof. Suppose $L \cap h = p$. Take $q \in L$ at distance 1 from p . Now take $r \in L$ distance $t > 1$ from p along the same direction as q . Let r' be a point on h of minimum distance to r .

Take the comparison triangle (P, R, R') and let Q' be the point on $[P, R']$ with $d(P, Q') = \frac{1}{t}d(P, R')$. The corresponding point q' lies in h along the geodesic from p to r' . Then we have

$$d(q, h) \leq d(q, q') \leq d(Q, Q') = \frac{1}{t}d(R, R') = \frac{1}{t}d(r, h).$$

Since $d(q, h)$ is a constant, we can make r arbitrarily far from h by increasing t . \square




FIGURE 2. The comparison triangle for Lemma 2.14

Here is an immediate consequence.

Corollary 2.15. *If L_1 and L_2 are axes of w , then L_1 properly intersects h if and only if L_2 does.*

Furthermore h has the structure of a cube complex, given by the midcubes it comprises. If X is CAT(0), it is straightforward to check that so is h . If w is a hyperbolic isometry of X preserving h , then w has an axis in h which must also be an axis in X .

2.6. Special cube complexes. The theory of special cube complexes was introduced by Haglund and Wise [HW08]. Special cube complexes are cube complexes whose hyperplanes avoid certain pathologies. They are defined in such a way that guarantees a locally isometric embedding in the Salvetti complex of some RAAG. Thus if C is a special cube complex, then there is an injective homomorphism from $\pi_1(C)$ to some RAAG.

Later, Hsu and Wise used special cube complexes to show that many free by cyclic graph groups virtually embed in RAAGs.

Theorem 2.16. [HW10] *Let G be a graph of groups with free vertex groups and cyclic edge groups. Suppose G contains no Baumslag-Solitar groups, or equivalently is word hyperbolic. Then G acts on a cube complex C , and G contains a finite index subgroup H whose action on C is special. As a result, H quasi-isometrically embeds in some RAAG.*

This theorem applies to a free by cyclic HNN extension that we'll use as an example later.

A final application is the following result, due to Ian Agol. It is part of his proof of the virtual Haken conjecture.

Theorem 2.17. [AGM12] *Let C be a non-positively curved cube complex such that $\pi_1(C)$ is word hyperbolic. Then C has a finite cover which is special.*

These two theorems are statements of existence. Neither gives a method for finding the desired cover.

3. $X^{(1)}$ GEODESICS

3.1. Interaction with hyperplanes. In this section, A will denote a RAAG. We'll use X to refer to the universal cover of its Salvetti complex.

Most of our arguments exploit our ability to relate the combinatorics of A to the geometry of X . We can, in fact, see the group directly in X by tracing our paths not in X but in its 1-skeleton.

We'll call paths in X that stay in $X^{(1)}$ and don't change direction mid-edge *edge paths*. Edge paths between vertices are nothing more than words in the generators of A . The path follows the edge corresponding to each letter in turn.

Unless A is free, $X^{(1)}$ is not CAT(0) and does not have unique geodesics. We'll use $[p, q]$ to denote some choice of minimal length path between p and q in $X^{(1)}$. We'll call this an $X^{(1)}$ *geodesic segment*. Given an edge path from p to q encoded by $w \in A$, finding an $X^{(1)}$ geodesic $[p, q]$ is equivalent to finding a shortest form for w . The distance from 1 to w in the word metric, denoted $|w|$ is also the length of $[p, q]$. Here are some quick results connecting $X^{(1)}$ geodesics segments to hyperplanes.

Observe that an edge path is transverse to the hyperplanes of X . Furthermore, two vertices of X lie on the same side of a hyperplane h if and only if every edge path between them crosses h an even number of times.

Lemma 3.1. *An $X^{(1)}$ -geodesic segment crosses no hyperplane more than once.*

Proof. Suppose $[p, q]$ does cross some hyperplane twice. Consider an innermost pair of crossings. Let h be the hyperplane they cross. The path between the crossings does not meet any hyperplane h' disjoint to h , otherwise it would have to recross h' in order to return to h , violating our choice of an innermost pair of crossings. Thus the path stays adjacent to h , in an $h \times I$ neighborhood. This means $[p, q]$ can be shortened by removing two crossings of h , so it is not a geodesic segment. \square

Corollary 3.2. *An $X^{(1)}$ geodesic segment is any segment that crosses no hyperplane more than once. Given two points, the X geodesic segment and $X^{(1)}$ geodesic segments between them all cross the same set of hyperplanes.*

The following corollary is an application of the triangle inequality.

Corollary 3.3. *Given $w \in A$ not equal to 1 and a vertex $p \in X$ such that $d(p, wp)$ is minimal, the set $\xi = \bigcup_{i \in \mathbb{Z}} w^i [p, wp]$ is an axis of w , that is a w -invariant bi-infinite*

$X^{(1)}$ geodesic.

Note that ξ depends both on the choice of p and $[p, wp]$, so it is far from unique. However, ξ does remain within bounded distance of any (non- $X^{(1)}$) axis L of w , since it is the w orbit of a (bounded) segment.

The last lemmas rely on the specific construction of X , rather than general CAT(0) hyperplane facts. By inspection, the hyperplanes of X/A meet a single edge each. Thus we can assign to each hyperplane in X a *type*, according to which hyperplane in X/A lifts to it. Moreover, since the hyperplanes of X/A do not self-intersect, distinct hyperplanes of the same type in X do not intersect. Finally, note that the action of A preserves type, as well as the orientation of the edges that cross each hyperplane.

Lemma 3.4 (Slope Lemma). *Suppose $w \in A$ has axis ξ . Suppose p is a point on ξ , and q is a point on a hyperplane h . h crosses ξ beyond wp (that is, on the ray $\xi = \bigcup_{i=1}^{\infty} w^i [p, wp]$) if and only if h crosses $[wp, wq]$. Note figure 3.*

Proof. If h crosses ξ then $wh \neq h$. If the crossing is beyond wp then h must cross $[wp, wq]$ to avoid crossing wh .




FIGURE 3. An illustration of the Slope Lemma

On the other hand, if h crosses $[wp, wq]$, we have that $w^{-1}h$ crosses $[p, q]$. Since $wh \neq h$, $w^{-1}h$ must also cross either ξ or $[wp, wq]$. Continuing by induction, we either get that some $w^{-k}h$ crosses ξ (thus producing the desired conclusion) or that infinitely many hyperplanes cross $[p, q]$, an absurdity. \square

If h_1 and h_2 are two hyperplanes in X whose projections in X/A intersect transversely we write $h_1 \perp h_2$. If their projections are disjoint or identical we write $h_1 \parallel h_2$. No hyperplane of X/A self intersects, so if $h_1 \parallel h_2$, then h_1 and h_2 are either identical or disjoint. However not all pairs $h_1 \perp h_2$ intersect.

If two hyperplanes $h_1 \perp h_2$ are adjacent to the same vertex, then they do in fact intersect in a square that meets that vertex. Also, only one edge of each type and direction meets a given vertex. These facts are immediate from our construction of X . They are also, however, the remaining two conditions for the action of A on X to be *special*. The following lemma is true for any group acting specially on a CAT(0) cube complex.

Lemma 3.5. *Let p be a vertex and h a hyperplane. The vertex p is adjacent to h if and only if every geodesic segment $[q, p]$ from any $q \in h$ crosses only hyperplanes $h_i \perp h$.*

Proof. Let p_i be the i^{th} vertex in $[q, p]$, and h_i be the hyperplane crossing $[p_i, p_{i+1}]$. Suppose $h_i \perp h$ for all i . We'll show p is adjacent to h by induction. The first vertex, p_1 , is adjacent by definition. Now suppose p_i is adjacent to h , via an edge we'll denote e . The edges $[p_i, p_{i+1}]$ and e span a square, h is a midcube of the square, and p_{i+1} is still adjacent to h .

On the other hand if some $h_i \parallel h$ then h_i separates h from $[p_{i+1}, p]$, since $[q, p]$ crosses no hyperplane twice. Thus any path from h to p must cross h_i , and p is not adjacent to h . \square

Any element of A that preserves h also preserves the set of adjacent vertices on each side. Given a point p adjacent to h , the stabilizer subgroup of h in A sends p along the edges that cross $h_i \perp h$. In a RAAG, these translations have a particular form. They correspond to the generators that commute with the generator whose edge crosses h . Thus the stabilizer subgroup is conjugate to a subgroup generated by vertices of Γ . This characterization allows us to conclude that if $w^i h = h$ then $wh = h$, which permits a useful variant of the Slope Lemma (3.4).

Lemma 3.6 (Parallel Axis Lemma). *Let $w \in A$ have axis ξ as above, and $p \in \xi$. Suppose that some hyperplane h does not meet ξ and r is some point in X such that $[p, r]$ meets h . If h intersects $w^n[p, r]$ for any $n \neq 0$, then $wh = h$.*




FIGURE 4. An illustration of the Parallel Axis Lemma

Proof. Let $q = h \cap [p, r]$. Suppose that for some $n \neq 0$, h meets $w^n[p, r]$ at q' . If $d(p, q) < d(w^n p, q')$, then the Slope Lemma implies that h meets ξ beyond p (on the ray not containing $w^n p$). Similarly, if $d(p, q) > d(w^n p, q')$, then h meets ξ beyond $w^n p$. However, h doesn't meet ξ at all. We conclude that $d(p, q) = d(w^n p, q')$ and $w^n q = q'$ (see Figure 4). Thus $w^n h = h$. As noted above, if w^n is in the stabilizer of h then so is w . \square

3.2. A standard form for $X^{(1)}$ geodesics. The Salvetti complex X/A has a single vertex. If we choose a base vertex of X , which we'll call 1, then every other vertex has the form $w \cdot 1$ for some $w \in A$. For brevity, we'll let w denote the vertex $w \cdot 1$ as well as the group element.

Given an element $w \in A$ we'll choose and label a few useful objects, also illustrated in Figure 5:

- (1) A vertex a_w on an axis of w that is minimal distance from 1 among all vertices that lie on axes of w
- (2) A geodesic \mathbf{w} from a to wa
- (3) A geodesic \mathbf{s}_w from 1 to a .
- (4) The concatenation $\overleftrightarrow{\mathbf{w}} = \bigcup_{i \in \mathbb{Z}} w^i \mathbf{w}$, which if $w \neq 1$ is an axis of w .




FIGURE 5. The geodesics \mathbf{s} , \mathbf{w} and ws . The translates of \mathbf{w} form an axis $\overleftrightarrow{\mathbf{w}}$.

Notice that if 1 lies on an axis of w , then $a_w = 1$. Also, our geodesics are in $X^{(1)}$, so these choices are not unique. Henceforth, we'll omit the subscripts when there is no ambiguity.

Lemma 3.7 (Standard Form). *Given an element w in A and a choice of \mathbf{w} , a and \mathbf{s} as above,*

(1) *No hyperplane meets $\mathbf{W} = \overleftrightarrow{\mathbf{w}} \cup \bigcup_{i \in \mathbb{Z}} w^i \mathbf{s}$ more than once.*

(2) *The concatenation of \mathbf{s} , \mathbf{w} and $w\mathbf{s}$ is a geodesic from 1 to w .*

Notation. Given a set of hyperplanes \mathcal{H} and an edge path γ , let \mathcal{H}_γ denote the set of hyperplanes in \mathcal{H} that meet γ an odd number of times. This is equivalent to saying \mathcal{H}_γ is the set of $h \in \mathcal{H}$ that separate the endpoints of γ , if it has them. Note $\mathcal{H}_{\gamma'} = \mathcal{H}_\gamma$ for any γ' with the same endpoints. If α, β, γ is a triangle, then $\mathcal{H}_\gamma = \mathcal{H}_\alpha \Delta \mathcal{H}_\beta$ (the symmetric difference).

Proof of Lemma 3.7. (1) Suppose some h meets \mathbf{W} more than once. We can assume (translating by w) that one of those intersections lies on \mathbf{s} , since no hyperplane meets an axis twice. We can further assume that h is adjacent to a . If it is separated from a by some other hyperplane, then that hyperplane also intersects \mathbf{W} more than once.

Let $p = h \cap \mathbf{s}$ and let q be a different point of $h \cap \mathbf{W}$. By Lemma 3.5, $[wa, wp]$ meets only hyperplanes of intersecting type with wh . Thus h does not meet $[wa, wp]$. Therefore, by the Slope Lemma (3.4), the point q cannot lie on $\overleftrightarrow{\mathbf{w}}$ beyond wa . Similarly q cannot lie on $\overleftrightarrow{\mathbf{w}}$ beyond $w^{-1}a$. This leaves the following possibilities for q :

- h meets $\overleftrightarrow{\mathbf{w}}$ between $w^{-1}a$ and wa .
- h does not meet $\overleftrightarrow{\mathbf{w}}$ at all. Thus $q \in w^n \mathbf{s}$ for some $n \neq 0$ and by the Parallel Axis Lemma (3.6), $wh = h$.

Thus, exploiting symmetry, we can take h to cross \mathbf{w} or we can stipulate $h = wh$ (but not both, since no hyperplane meets $\overleftrightarrow{\mathbf{w}}$ twice). In either case, we reach the same contradiction as follows:

Let a' be the vertex adjacent to a that shares the edge through h . Let \mathcal{H} be the set of all hyperplanes of X . Then we have $\mathcal{H}_{[1, a']} = \mathcal{H}_\mathbf{s} - \{h\}$. We also have $\mathcal{H}_{[a', wa']} = \{h\} \Delta \mathcal{H}_\mathbf{w} \Delta \{wh\}$, which in either case means that $d(a', wa') = d(a, wa)$. We conclude a' lies on an axis and is closer to 1 than a , violating our construction of a .

(2) Since the concatenation of \mathbf{s} , \mathbf{w} and $w\mathbf{s}$ crosses every hyperplane between 1 and w once, but none of them twice, it must be a geodesic segment. \square

Remark. The use of Lemma 3.5 is a convenience here. One can prove this lemma with only the assumption that X is CAT(0), and hyperplanes of X/A don't self-intersect.

Lemma 3.8 (Separating Lemma). *Let $w \neq 1$ be an element of A and suppose that a hyperplane h is disjoint from one (and hence every) $X^{(1)}$ axis of w . Suppose that $wh \neq h$. Then there is a hyperplane $h_1 \parallel h$ that separates h from wh . h_1 does not intersect $w^i h$ for any $i \in \mathbb{Z}$.*

Proof. Let $[p, a]$ be an edge path from h to an axis of w , minimal in length among all such paths for all axes of w . p is not a vertex, so we'll name the vertices of the edge it lies on. Let q be the one that lies in $[p, a]$. Set the other vertex as the basepoint 1. Then $[1, w \cdot 1]$ decomposes into a geodesics \mathbf{s} , \mathbf{w} and $w\mathbf{s}$, with

$\mathbf{s} = [1, a] \supset [q, a]$. By the Standard Form Lemma (3.7),

$$\begin{aligned} h \cap [1, w] &= \{p\} \\ wh \cap [1, w] &= \{wp\} \\ w^i h \cap [1, w] &= \emptyset \text{ for } i \neq 0, 1. \end{aligned}$$

We conclude that $[q, wq]$ meets no $w^i h$.

Suppose that $[q, wq]$ meets no hyperplanes $h_1 \parallel h$. Then by Lemma 3.5, wq is adjacent to both h and wh . But w preserves orientation, and only one edge of each type and direction meets the vertex wq . This implies that $wh = h$, which violates our hypothesis. We conclude that $[q, wq]$ does cross a hyperplane $h_1 \parallel h$, which must therefore separate h from wh . For all $i \notin \{0, 1\}$, the hyperplane $w^i h \parallel h_1$. Notice though, that h_1 meets $[q, wq]$ and $w^i h$ doesn't. Thus $w^i h$ can't be equal to h_1 , so they are disjoint. \square

Repeated application of this lemma gives a corollary.

Corollary 3.9. *As in the lemma let $w \neq 1$ be an element of A and suppose that a hyperplane h_0 is disjoint from one (and hence every) $X^{(1)}$ axis of w . Suppose that $wh_0 \neq h_0$. Then there is a hyperplane h_N meeting $\overleftarrow{\mathbf{w}}$ that separates h_0 from wh_0 , and does not intersect $w^i h_0$ for any $i \in \mathbb{Z}$.*

Proof. We'll argue by induction. Given an h_n that satisfies the hypotheses of Lemma 3.8 we produce h_{n+1} . If h_{n+1} crosses $\overleftarrow{\mathbf{w}}$, then we are done. If not, then it crosses \mathbf{s} or $w\mathbf{s}$, but not both. This means that h_{n+1} separates either h_n or wh_n from $\overleftarrow{\mathbf{w}}$. Furthermore, $wh_n \neq h_n$, and we can apply the Lemma again. Since \mathbf{s} crosses finitely many hyperplanes, repeated application of the lemma will eventually produce some h_N that crosses $\overleftarrow{\mathbf{w}}$. Each hyperplane $w^i h_0$ is separated from h_N by a sequence of hyperplanes $w^i h_n$. \square

4. TWO-GENERATOR SUBGROUPS

4.1. Essential hyperplanes. A two-generator subgroup of a RAAG A is the image of a homomorphism $\phi : F \rightarrow A$, where F is a free group on two generators. A acts on X , the universal cover of its Salvetti complex, so F acts on X too. We'll produce the following tree T containing F and produce a map $\phi : T \rightarrow X$ which extends a standard orbit map on F .

First we assume that $F = \langle u, v \rangle$ and we choose a basepoint 1 on some axis of u . We then choose \mathbf{u} , a_v , \mathbf{s}_v , \mathbf{v} as defined in the previous section (note that $a_u = 1$, so we don't bother with \mathbf{s}_u), and immediately drop the subscripts.

We define T to be a trivalent tree with vertices labeled w and wa for all $w \in F$. The edges of T are as follows:

- (1) w is adjacent to wu via edge $w\mathbf{u}$
- (2) w is adjacent to wa via edge $w\mathbf{s}$
- (3) wa is adjacent to wva via edge $w\mathbf{v}$

We extend the map ϕ to T by mapping each vertex to its eponymous vertex in X and each edge to its eponymous geodesic in some equivariant way (say, constant speed). This has the effect that for some $q \in T$ with $\phi(q) \in h$, we have $\phi(wq) \in h$ if and only if $wh = h$.

Definition 4.1. We say a hyperplane h is *essential* if $\phi^{-1}(h)$ is a single point.




FIGURE 6. A diagram of T near the vertex 1

The advantage of essential hyperplanes to a metric argument should be clear. A reduced word $w \in F$ can be taken as a geodesic in T . If it crosses $\phi^{-1}(h)$ once, and h is essential, then it will never cross $\phi^{-1}(h)$ again. Thus our strategy is to produce sufficiently many of these to justify the main theorem. Our tool for proving that a hyperplane is essential is the following definition and lemma.

Definition 4.2. Suppose $p \in T$. We denote T_p to be the closure of the connected component of $T - \{wp \mid w \in F - \{1\}\}$ containing p . See Figure 7.




FIGURE 7. Diagrams of T_p for two possible p .

Lemma 4.3. *If $\phi^{-1}(h) \cap T_p = \{p\}$, then h is essential.*

Proof. There are three types of edges in T , and within each type, the preimages of hyperplanes are identically distributed (and finite). Thus if h is not essential, we can pick a point p' of $\phi^{-1}(h) \setminus \{p\}$ whose distance to p in T is minimal.

Let γ be the geodesic in T from p to p' . Since $p' \notin T_p$ by hypothesis, γ contains some point $q = wp$ in the boundary of T_p . By assumption, $q \notin h$ so $wh \neq h$. Since p and p' lie of the same side of wh , we conclude that γ must cross wh again at some q' , either between p and q or between q and p' . Translating by w^{-1} we have that $w^{-1}q'$ lies in $\phi^{-1}(h)$ and is closer to p than p' , contradicting our choice of p' . \square

4.2. The existence of essential hyperplanes. The following proposition comprises much of the work in proving Theorem 1.4. Recall the following notation: Given a set of hyperplanes \mathcal{H} , and a geodesic γ , \mathcal{H}_γ denotes the set of hyperplanes of \mathcal{H} that intersect γ .

Proposition 4.4. *If $\phi(F)$ is not abelian, then we can choose a basepoint 1 and generators u, v of F such that for the associated T , there exists an essential hyperplane h in X meeting $[1, v]$.*

Proof. Let $\{u_{\text{orig}}, v_{\text{orig}}\}$ be a generating set for F . Pick a hyperplane that crosses an axis of the commutator $c_{\text{orig}} = u_{\text{orig}}v_{\text{orig}}u_{\text{orig}}^{-1}v_{\text{orig}}^{-1}$. Let \mathcal{H} be the set of all hyperplanes of that type.

Now pick a new generating set u, v and a basepoint 1 such that the following triple of numbers is minimal (lexicographically) among all possible choices of u, v and 1:

$$(\#\mathcal{H}_u, \#\mathcal{H}_{[1,v]}, \#\mathcal{H}_v).$$

We claim $\mathcal{H}_{[1,v]}$ contains an essential hyperplane. There are two cases which, unfortunately, admit almost no overlap in their discussion.

Case 1: Suppose \mathcal{H}_u is nonempty. We will argue that each of $\mathcal{H}_u, \mathcal{H}_v$ contains an essential hyperplane.

Since the hyperplanes of \mathcal{H} are pairwise disjoint, we can order the \mathcal{H}_u by distance from 1, and let h_1, h_2 be the middle ones (if $\#\mathcal{H}_u$ is odd, $h_1 = h_2$). We claim that one of h_1, h_2 does not meet the nearby v -paths, seen also in Figure (8a):

$$[1, v], [1, v^{-1}], [u, uv], [u, uv^{-1}].$$




FIGURE 8. h_i and the nearby v -paths

Similarly if we know that \mathcal{H}_v is nonempty, otherwise we could switch u and v and take basepoint a to reduce the triple. Take h'_1 and h'_2 to be the middle hyperplanes of \mathcal{H}_v . Then these are also the middle hyperplanes of $\mathcal{H}_{[1,v]}$, since $\#\mathcal{H}_s = \#\mathcal{H}_{v_s}$. We make an analogous claim that one of h'_1 and h'_2 meets none of the nearby u -paths. The following three facts establish our claims about the h_i and h'_i ($i \in \{1, 2\}$).

- (1) Neither h_i meets either of the “far” v -paths. For instance, if h_1 met $[u, uv^{-1}]$, then $[1, uv^{-1}]$ would meet fewer hyperplanes than $[u, uv^{-1}]$, as seen in Figure 8b. Defining $v_{\text{new}} = uv^{-1}$ would give

$$\#\mathcal{H}_{[1, v_{\text{new}}]} < \#\mathcal{H}_{[u, uv^{-1}]} = \#\mathcal{H}_{[1, v]}.$$

Thus $\{u, v_{\text{new}}\}$ would produce a smaller triple than $\{u, v\}$, violating our minimality assumption. Similarly, neither h'_i meets either of the far u -paths.

In the case that $h_1 = h_2$, all four v -paths are “far,” so the claim is proven.

- (2) The h_i do not each meet the near v -paths on the same side. For instance, suppose h_1 met $[1, v]$. We claim that h_2 does not meet $[u, uv]$. If it did, then $h_2 = uh_1$ as in Figure 9a, but no hyperplane meets $\overleftarrow{\mathbf{u}}$ twice. Similarly, since h'_1 and h'_2 meet the axis of v , we have h'_1 and h'_2 do not both meet the near u -paths on the same side.
- (3) The h_i do not each meet the near v -paths on the opposite sides. For instance, suppose h_1 met $[1, v]$ and h_2 met $[u, uv^{-1}]$.

Choose a point $e \in \mathbf{u}$ between h_1 and h_2 . Notice $\mathcal{H}_{[1, e]}$ consists of h_1 and all the hyperplanes of \mathcal{H} that separate it from 1, while $\mathcal{H}_{[e, u]}$ consists of h_2 and all the hyperplanes of \mathcal{H} that separate it from u .

By part (1), h_1 does not meet the far v -path $[u, uv^{-1}]$, so it meets $[1, vu^{-1}]$. Also by part (1) no hyperplane of $\mathcal{H}_{[1, v]}$ meets both $[1, u]$ and $[v, vu^{-1}]$ (one of these u -paths would be “far”). Since $h_1 \in \mathcal{H}_{[1, v]}$, it meets $[u, uv^{-1}]$ (see Figure 9b).

Thus if we set $v_{\text{new}} = uv^{-1}$ we have that h_1 meets all of $[1, u]$, $[1, v_{\text{new}}]$ and $[1, v_{\text{new}}^{-1}]$. So $\mathcal{H}_{[1, e]}$ is a subset of $\mathcal{H}_{\mathbf{u}}$, $\mathcal{H}_{[1, v_{\text{new}}]}$, and $\mathcal{H}_{[1, v_{\text{new}}^{-1}]}$. Translating the last inclusion by v_{new} gives $\mathcal{H}_{[v_{\text{new}}, v_{\text{new}}e]} \subset \mathcal{H}_{[1, v_{\text{new}}]}$. Let's see what this implies.

$$\begin{aligned} \mathcal{H}_{[e, ue]} &= \mathcal{H}_{[1, e]} \Delta \mathcal{H}_{\mathbf{u}} \Delta \mathcal{H}_{[u, ue]} = (\mathcal{H}_{\mathbf{u}} - \mathcal{H}_{[1, e]}) \cup \Delta \mathcal{H}_{[u, ue]} \\ \text{so } \#\mathcal{H}_{[e, ue]} &= \#\mathcal{H}_{\mathbf{u}}. \end{aligned}$$

$$\begin{aligned} \mathcal{H}_{[e, v_{\text{new}}e]} &= \mathcal{H}_{[1, e]} \Delta \mathcal{H}_{[1, v_{\text{new}}]} \Delta \mathcal{H}_{[v_{\text{new}}, v_{\text{new}}e]} \\ &= \mathcal{H}_{[1, v_{\text{new}}]} - (\mathcal{H}_{[1, e]} \Delta \mathcal{H}_{[v_{\text{new}}, v_{\text{new}}e]}) \subsetneq \mathcal{H}_{[1, v_{\text{new}}]} \\ \text{so } \#\mathcal{H}_{[e, v_{\text{new}}e]} &< \#\mathcal{H}_{[1, v_{\text{new}}]}. \end{aligned}$$

Out hypothesis on h_2 gives $\mathcal{H}_{[e, u]} = \mathcal{H}_u \cap \mathcal{H}_{[u, uv]}$, so $\mathcal{H}_{[1, v_{\text{new}}]} = \mathcal{H}_u \Delta \mathcal{H}_{[u, uv]} = \mathcal{H}_{[1, e]} \sqcup (\mathcal{H}_{[u, uv]} - \mathcal{H}_{[e, u]})$. So $\#\mathcal{H}_{[1, v_{\text{new}}]} = \#\mathcal{H}_{[1, v]}$.

This means that replacing v by v_{new} and shifting the basepoint to e would strictly reduce the triple, violating our assumptions about its minimality under u , v and 1. The v case is similar.

This proves our claim about h_1 and h_2 in $\mathcal{H}_{\mathbf{u}}$. Call one that meets none of the nearby v -paths h_u . Similarly call one of h'_1, h'_2 that meets none of the nearby u -paths h_v . Furthermore, h_v is separated from each edge $[v^n, v^n u^{\pm 1}]_{n \notin \{0, 1\}}$ by one of its own v -translates.

We'll use Lemma 4.3 to show that h_u is essential. Consider $p = h_u \cap \mathbf{u}$. Then to see that $T_p \cap h_u = \{p\}$ we check:

- (1) $\overleftarrow{\mathbf{u}} \cap h_u = \{p\}$




FIGURE 9. Illustration of facts (2) and (3)

- (2) h_u doesn't meet the four nearby v -paths: $[1, v], [1, v^{-1}], [u, uv], [u, uv^{-1}]$ by our choice of h_u .
- (3) h_u does not meet any point that lies across any of the hyperplanes: $h_v, uh_v, v^{-1}h_v$, or $uv^{-1}h_v$ (see Figure 10), since none of these cross h_u .




FIGURE 10. Verifying h_u is essential. Notice h_u doesn't cross these 4 translates of h_v .

This verifies that $T_p \cap h_u = \{p\}$ and by Lemma 4.3, h_u is essential. By identical reasoning, so is h_v . As it lies in $\mathcal{H}_{[1,v]}$, h_v is the hyperplane we sought.

Case 2: Suppose \mathcal{H}_u is empty. Recall that \mathcal{H} is the set of all hyperplanes of some type that crosses the axis of c_{orig} , the commutator of the original generators of F . However, one can check that $c = uvu^{-1}v^{-1}$ is equal to $gc_{\text{orig}}^{\pm 1}g^{-1}$ for some $g \in A$ (the reader can check this on his or her favorite generating set of $\text{Aut}(F)$). Thus \mathcal{H} also contains a hyperplane that crosses any axis of c . Consider a standard decomposition of $[1, c]$ into s_c, c and cs_c . Since \mathcal{H} is c -invariant, there is some $h \in \mathcal{H}$ that meets $c^{-1}c$. Our goal is to show that this h is essential.

Consider the geodesic from 1 to c^{-1} in T . It maps to a piecewise geodesic path ξ in X , which is a concatenation of two translates of $[1, v]$ and two translates of u . Since \mathcal{H}_u is empty, h meets ξ at either $[1, v]$ or $[vu, vuv^{-1}]$, but not both (see Figure 11). Assume, without loss of generality, that it meets $[1, v]$ at p .

We produce two facts about the stabilizer of h . Since $vuv^{-1}p \notin h$, we know $vuv^{-1}h \neq h$. We also claim that $uh \neq h$. If we suppose otherwise, then h meets $c\xi$




FIGURE 11. The intersection of $\phi^{-1}(h)$ with ξ . If $h \in \mathcal{H}_{[1,v]}$ intersects $c^{-1}\mathbf{c}$ and $\mathcal{H}_{\mathbf{u}}$ is empty, then $\phi^{-1}(h)$ meets ξ exactly once and doesn't meet $c^{-1}\xi$ at all.

at up . Since h cannot cross $[1, c]$ or any translate of \mathbf{u} , it must meet $c\xi$ a second time, and that meeting must be some $q \in [uvu^{-1}, c]$. Since h doesn't meet $uv\overleftarrow{\mathbf{u}}$, the Parallel Axis Lemma says that $q = cp$. But if $ch = h$, then h cannot meet $\overleftarrow{\mathbf{c}}$ at all, violating our choice of h . Thus $uh \neq h$ as claimed. We'll use the fact that neither uh nor $vuv^{-1}h$ is equal to h in the argument that follows.

We'll now argue that h is essential using Lemma 4.3. To do this we will consider the intersection of h with the following 5 sets (illustrated in Figure 12). The preimages of these sets cover T_p (where p now refers to $\phi^{-1}(h) \cap [1, v] \in T$) regardless of whether p lies on \mathbf{s} , \mathbf{v} or vs . Computing these intersections will show that $\phi^{-1}(h) \cap T_p = \{p\}$.

- (1) $\mathbf{V} = \overleftarrow{\mathbf{v}} \cup \bigcup_{n \in \mathbb{Z}} v^n \mathbf{s}$. h meets \mathbf{V} only at p , by Lemma 3.7.
- (2) $\overleftarrow{\mathbf{u}} \cup v \overleftarrow{\mathbf{u}}$. The intersection of h with these axes is empty, because no F -translate of h meets \mathbf{u} .
- (3) $\bigcup_{n \neq 0} [u^n, u^n v]$. Since h doesn't meet $\overleftarrow{\mathbf{u}}$, the Parallel Axis Lemma (3.6) states that if h meets one of these edges then $uh = h$. We've shown that this isn't the case, so h does not intersect this set.
- (4) $\bigcup_{n \neq 0} [vu^n v^{-1}, vu^n]$. Note as in the previous step that h doesn't meet $v \overleftarrow{\mathbf{u}}$ and $vuv^{-1}h \neq h$. Thus h does not intersect this set.
- (5) Finally, in the case $p \in \mathbf{v}$, the region T_p overlaps the paths $[u^n v^{-1}, u^n]$ and $[vu^n, vu^n v]$. We'll use the following v -**straightening** argument to show that h doesn't meet $[u^n v^{-1}, u^n]$, noting that a similar one exists for $[vu^n, vu^n v]$:

Suppose h meets $[u^n v^{-1}, u^n]$. Since $\mathcal{H}_{\mathbf{u}}$ is empty, $[1, v]$ meets the same hyperplanes as $[1, vu^{-n}]$. So if we set $v_{\text{new}} = vu^{-n}$, then $\mathcal{H}_{[1, v_{\text{new}}]} = \mathcal{H}_{[1, v]}$. We may then produce a T_{new} with \mathbf{v}_{new} and \mathbf{s}_{new} . But h and every hyperplane of \mathcal{H} between 1 and h (including those of $\mathcal{H}_{\mathbf{s}}$), meet both $[1, v_{\text{new}}]$ and $[1, v_{\text{new}}^{-1}]$. Thus $h \cup \mathcal{H}_{\mathbf{s}} \subset \mathcal{H}_{\mathbf{s}_{\text{new}}}$, which means $\#\mathcal{H}_{\mathbf{v}_{\text{new}}} < \#\mathcal{H}_{\mathbf{v}}$, violating the minimality of the triple.

This establishes that $\phi^{-1}(h) \cap T_p = \{p\}$. By Lemma 4.3, h is essential. \square

The following corollary may be of interest, but we won't use it here.

Corollary 4.5. *For $F \rightarrow A$ a homomorphism and any hyperplane h , there is a choice of generators $u, v \in F$ and a basepoint 1 such that either*

- (1) h does not intersect T .




FIGURE 12. T_p is divided into 5 regions to verify that $h \cap T_p = \{p\}$. This is the $p \in \mathbf{v}$ case.

- (2) There exists $w \in F$ such that wh is essential at some $p \in [1, v]$.
(3) There exists $w \in F$ such that wh meets $[1, v]$ and either $uh = h$ or $vwv^{-1}h = h$ (possibly both).

Proof. Apply the methods of the previous proposition to $\mathcal{H} = Fh$. Assuming the conditions of case 1 leads to essential hyperplanes $h_u \in \mathcal{H}_{\mathbf{u}}$ and $h_v \in \mathcal{H}_{\mathbf{v}}$. Since $h_u = wh_v$ for some $w \in F$, we have that h_v meets $w^{-1}\mathbf{u}$. This contradicts the claim that it is essential. In case 2 we showed that every hyperplane of $\mathcal{H}_{[1, v]}$ is essential unless $uh = h$ or $vwv^{-1}h = h$. So either $\mathcal{H}_{[1, v]}$ is empty (along with $\mathcal{H}_{\mathbf{u}}$ by assumption) or some w -translate of h satisfies (2) or (3). \square

4.3. The main theorem. Using Proposition 4.4, we are ready to prove the main theorem: that ϕ is a quasi-isometric embedding.

Proof of Theorem 1.4. For a word w of length n , we know that

$$|\phi(w)| \leq n \max\{|\phi(u)|, |\phi(v)|\}.$$

It remains to find a lower bound that is linear in n . We'll produce a pair of essential hyperplanes, meeting \mathbf{u} and \mathbf{v} . We'll consider their orbits Fh_u and Fh_v , and count how many of these hyperplanes cross $[1, w]$. Note that in case 1 of Proposition 4.4 we already produced an adequate pair of essential hyperplanes, but the methods here only rely on the assumption that some essential h crosses $[1, v]$.

Take the choice of u, v and a basepoint from Proposition 4.4. Consider all possible generators $v_{\text{new}} = u^{n_1}vu^{n_2}$, and the T_{new} generated by $\langle u, v_{\text{new}} \rangle$, but keeping the same basepoint. Let h be the hyperplane of X that lies closest to $\overleftarrow{\mathbf{u}}$ among all essential (with respect to T_{new}) hyperplanes meeting $[1, v_{\text{new}}]$. Such hyperplanes exist by Proposition 4.4, and a closest one exists because $\overleftarrow{\mathbf{u}}$ is the u orbit of the compact set \mathbf{u} . Let $p = \phi^{-1}(h)$.

The Separating Lemma (3.8) part (1) states that there is a hyperplane $h' \parallel h$ between h and uh . Since any F -translate h meets T only once, we know $h' \notin Fh$ so it is disjoint from all F -translates of h . As a result, $h' \cap T$ is contained in T_p . In fact, it is contained in the component of $T_p - \{p\}$ that contains uh .

We claim h' crosses $\overleftarrow{\mathbf{u}}$. Suppose it does not. We will argue that h' satisfies the conditions that lead us to choose h but lies closer to $\overleftarrow{\mathbf{u}}$. h' meets $[p, up]$, so h' meets either $[1, p]$ or $[u, up]$ but not both. This means that $uh' \neq h'$. Without loss of generality, suppose it meets $[1, p]$. The contrapositive of the Parallel Axis

Lemma (3.6) therefore applies, meaning h' does not meet any $[u^n, u^n p]$ for $n \neq 0$. The only other possible point of $\phi^{-1}(h')$ is on some $[u^n v^{-1}, u^n]$. If h' doesn't meet any such edge, then it is essential and closer to $\overleftarrow{\mathbf{u}}$ than h . This contradicts our choice of h . If h' does meet this edge, then set $v_{\text{new}} = vu^{-n}$, and one may check that $T_{\text{new}} \cap \phi^{-1}(h')$ is a single point on \mathbf{s}_{new} . Thus h' would be essential for this v_{new} and closer to $\overleftarrow{\mathbf{u}}$ than h . This also contradicts our choice of h .

So h' meets some $u^m \mathbf{u}$. Now we claim that for some v_{new} , both h and h' are essential on T_{new} . That h is essential is immediate. If $v_{\text{new}} = u^{n_1} v u^{n_2}$, then $[1, v_{\text{new}}]$ meets $u^{n_1} h$ and no other F -translates. Thus $u^{n_1} h$ is essential on $[1, v_{\text{new}}]$.

As noted above, $h' \cap T$ is limited to one component of $T_p - \{p\}$. If h' meets any u -translate of $[1, p]$, then let $u^{k_1}[1, p]$ be the one farthest from $u^m \mathbf{u}$. Repeated applications of the Slope Lemma show that only finitely many translates of $[1, p]$ meet h' . They also show that h' meets every translate of $[1, p]$ between $u^{k_1}[1, p]$ and $u^m \mathbf{u}$. We define $u^{k_2}[1, v^{-1}p]$ similarly to produce the farthest translate meeting h' , and apply the Slope Lemma similarly. Then we can choose $v_{\text{new}} = u^{\sigma(m, k_1)} v u^{-\sigma(m, k_2)}$ where

$$\sigma(m, k) = \begin{cases} k - m - 1 & \text{if } k \leq m \\ k - m & \text{if } k > m \\ 0 & \text{if } k \text{ is not defined} \end{cases}$$

One can check that h' is essential, meeting T_{new} only once, on $u^m \mathbf{u}$.




FIGURE 13. An example of a non-essential h' meeting $u^2 \mathbf{u}$. The set $h' \cap T$ is limited to one component of $T_p - \{p\}$. In this case $k_1 = 1$ and $k_2 = 3$.

Now let $h_u = u^{-m} h'$. As noted above $u^{\sigma(m, k_1)} h$ is essential. It lies on either \mathbf{v}_{new} , \mathbf{s}_{new} or $v_{\text{new}} \mathbf{s}_{\text{new}}$. In the first case, we'll call it h_v .

In the second two cases, Corollary 3.9 produces a h_v meeting \mathbf{v}_{new} that doesn't intersect $u^{\sigma(m, k_1)} h$ or any $u^{\sigma(m, k_1)} v_{\text{new}}^n h$. This restriction, along with the Standard Form Lemma (3.7), means that h_v is essential.

Either way we now have essential hyperplanes h_u and h_v . The group F acts freely on the orbits Fh_u and Fh_v , which consist entirely of essential hyperplanes. Any geodesic in T_{new} which crosses the preimage $\phi^{-1}(h)$ of an essential hyperplane h will never cross that preimage again. Since we have one such hyperplane in $Fh_u \cup Fh_v$ for each translate of \mathbf{u} and \mathbf{v}_{new} , a reduced word w of length n in $\{u, v_{\text{new}}\}$ will cross exactly n hyperplanes of $Fh_u \cup Fh_v$.

We have been flexible with our choice of generators. Given a different set, say $\{u_{\text{orig}}, v_{\text{orig}}\}$, we might have to compose with some automorphism of F . But automorphisms are quasi-isometries, and quasi-isometric embeddings are closed under composition. \square

Remark. Our procedure is nearly algorithmic. While some steps assume that a certain quantity is minimized, we also showed explicitly how to reduce it, should the properties we required be absent. Computationally one needs to be able to check whether some wh is equal to h and whether some $[p, q]$ crosses h . One also must be able to produce the hyperplanes of the Separating Lemma. Neither of these tools is more difficult than placing an element of A_Γ in a standard form. Given those computational tools, the initial description of ϕ , and the steps of the proof, one can produce $1, u, v$, and the essential h_u, h_v .

The non-abelian case of Baudisch's Theorem is recovered by the following corollary. The results listed in the introduction follow quickly.

Corollary 4.6. *If $\phi(F)$ is not abelian, then ϕ is an injection.*

Proof. By the previous theorem, given any choice of generators (and hence word metric), ϕ is a quasi isometric embedding. There is a $\lambda \geq 1$ and $\epsilon \geq 0$ such that for all $w \in F$,

$$\frac{1}{\lambda}|w| - \epsilon \leq |\phi(w)|.$$

Suppose $w \in F$ is not the identity. It translates along its axis in T by some distance d . For sufficiently large n , we have $|w^n| \geq nd > \lambda\epsilon$. But then we have

$$0 = \frac{1}{\lambda}\lambda\epsilon - \epsilon < \frac{1}{\lambda}|w^n| - \epsilon \leq |\phi(w^n)|.$$


Since $\phi(w^n)$ is not the identity, neither is $\phi(w)$. \square

Proof of Theorem 1.2. Let H be a two-generator subgroup of A_Γ . RAAGs are torsion free [Bau81, 2.3], so if H is abelian, it is free abelian. If H is not abelian, then apply the previous lemma. \square

Proof of Corollary 1.5. As above, suppose first H is abelian. Free abelian subgroups of isometries of X are quasi-isometrically embedded in X (and hence in A_Γ) by the Flat Torus Theorem (presented in [BH99, II.7.1 and II.7.17]). If H is not abelian then Theorem 1.4 applies. \square

5. COUNTEREXAMPLES IN OTHER SETTINGS

5.1. A three-generator subgroup. The result of Baudisch cannot be extended to subgroups of more generators. He produces a counterexample himself [Bau81, 6.2], which we repeat here. Let Γ be the following graph:



Write $F_3 = \langle u, v, w \rangle$ and let $\phi : F_3 \rightarrow A_\Gamma$ be defined by

$$\phi(u) = ax \quad \phi(v) = by \quad \phi(w) = cz$$

Notice that no two of $\phi(u), \phi(v), \phi(w)$ commute, yet Baudisch shows that the image of ϕ is not a free group, meaning that ϕ is not an injection. He goes farther and shows that the image of ϕ is not even a RAAG.

5.2. An obstruction to embedding in a RAAG. Consider the group $\pi_1(M)$, where M is the figure-8 knot complement. $\pi_1(M)$ can be realized as a semi-direct product $F_2 \rtimes \mathbb{Z}$ with the following presentation.

$$\pi_1(M) = \langle a, b, t \mid tat^{-1} = ab, tbt^{-1} = bab \rangle$$

Now let $F_2 = \langle a, b \rangle$ and $i : F_2 \rightarrow \pi_1(M)$ be the inclusion. The subgroup F_2 is free.

Consider the element $t^n a t^{-n}$, which has length at most $2n + 1$ in $\pi_1(M)$. We can equate this to a word in a, b . We use the relations to cancel a power of t and t^{-1} and repeat n times. Each cancellation at least doubles the number of a 's and b 's. No a 's or b 's ever cancel, since we only produce positive powers of a and b . The result is that $t^n a t^{-n}$ is equal to a word of length at least 2^n in a and b . Thus there is an infinite sequence of elements $w_n \in F_2$ such that $|i(w_n)| \leq 2 \log_2 |w_n| + 1$. Thus i is not a quasi-isometric embedding. We can thus apply the following corollary to $\pi_1(M)$.

Corollary 5.1. *Let G be any finitely-generated group with a free two-generator subgroup $F_2 = \langle a, b \rangle$ that is not quasi-isometrically embedded. No RAAG has a subgroup isomorphic to G .*

Proof. Consider the inclusion $i : F_2 \rightarrow G$. As noted at the beginning of the proof of Theorem 1.4, for all $w \in F_2$, we have an upper bound $|i(w)| \leq |w| \max\{|i(a)|, |i(b)|\}$. Thus in order for i to fail to be a quasi-isometric embedding, F_2 must contain an infinite sequence of elements w_n such that $|i(w_n)|$ is sublinear with respect to $|w_n|$.

Let $\{g_i\}$ be a generating set for G . If there were an injective homomorphism $\alpha : G \rightarrow A_\Gamma$, then $\alpha \circ i$ would be an injection. For all w_n we would have $|\alpha(w_n)| \leq |i(w_n)| \max\{|\alpha(g_i)|\}$, which would still be sublinear with respect to $|w_n|$. Thus the map $\alpha \circ i$ violates Theorem 1.4. \square

Note however that, while $\pi_1(M)$ is not isomorphic to a subgroup of any RAAG, it is a hyperbolic free by cyclic HNN extension. The theorem of Hsu and Wise [HW10] (reproduced as Theorem 2.16 here) applies. Thus $\pi_1(M)$ has a finite index subgroup H that is isomorphic to a subgroup of some RAAG. It would be interesting to identify such an H . Here is a first step.

Proposition 5.2. *If H is a finite index normal subgroup of $\pi_1(M)$ with abelian quotient, then H is not isomorphic to a subgroup of any RAAG.*

Proof. Let H be an index m normal subgroup of $\pi_1(M)$. If $\pi_1(M)/H$ is abelian, then the commutator subgroup, $[\pi_1(M), \pi_1(M)]$, is a subgroup of H . So $a^{-1}tat^{-1} = b \in H$. Also, $tbt^{-1}b^{-1} = ba \in H$, so $a \in H$. Therefore, $F_2 \subset H$. Finally, $t^m \in H$. Thus we have a sequence of $w_n \in F_2$ such that $i(w_n) = t^{mn} a t^{-mn}$ and $|i(w_n)| \leq 2 \log_2 |w_n| + 1$. Thus H is not isomorphic to a subgroup of any RAAG by Corollary 5.1. \square

Hsu and Wise’s construction in [HW10] gives a cube complex C with an action by $\pi_1(M)$. The proposition above shows that the action of H is not special. Equivalently, no abelian cover of $C/\pi_i(M)$ is special.

Question. What is the smallest index subgroup of $\pi_1(M)$ that embeds in a RAAG? Does this number change for different hyperbolic groups of the form $F_2 \rtimes \mathbb{Z}$?

Question. Can we find groups acting on CAT(0) cube complexes such that no subgroup of a given index acts has a special action? Can we find these in the class of 3-manifold groups?

5.3. A badly embedded free subgroup. Baudisch’s three-generator subgroup shows that the most naive generalization of our main theorem fails to hold when we increase the number of generators of F . However, even assuming that F_n injects into A_Γ is insufficient to guarantee a quasi-isometric embedding.

As noted before, if M is the figure-8 knot complement, then $\pi_1(M)$ does not embed in any RAAG, but some finite index subgroup $H < \pi_1(M)$ does. The group $H \cap F_2$ is a finite index subgroup of F_2 and thus finitely generated [Sch27], [LS77, 3.9]. Call it F_n . The inclusions of H into $\pi_1(M)$ and F_n into F_2 are quasi-isometries (this is a standard exercise).

$$\begin{array}{ccccc} F_n & \xrightarrow{i} & H & \xrightarrow[\text{q.i.}]{\alpha} & A_\Gamma \\ i_F \downarrow & & \downarrow i_{\pi_1(M)} & & \\ F_2 & \xrightarrow{i} & \pi_1(M) & & \end{array}$$

It is straightforward to show that when a quasi-isometric embedding is composed with a map that is not a quasi-isometric embedding, the result is not a quasi-isometric embedding. We conclude that $\alpha \circ i$ is an injective homomorphism from free group to a RAAG that is not a quasi-isometric embedding.

Question. What is the minimum rank n such that some RAAG has a free subgroup of rank n whose inclusion is not a quasi-isometric embedding?

REFERENCES

- [AB90] S. B Alexander and R. L Bishop. The hadamard-cartan theorem in locally convex spaces. *Enseign Math.*, 36:309–320, 1990.
- [AGM12] Ian Agol, Daniel Groves, and Jason Manning. The virtual Haken conjecture. *Preprint*, 2012. arXiv:1204.2810 [math.GT].
- [Bau81] A. Baudisch. Subgraphs of semifree groups. *Acta Mathematica Academiae Scientiarum Hungaricae*, 38(1-4):19–28, 1981.
- [BH99] Martin Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*. Springer, 1999.
- [CD95] Ruth Charney and Mike Davis. Finite $k(\pi, 1)$ s for Artin groups. In Frank Quinn, editor, *Prospects in Topology: Proceedings of a Conference in Honor of William Browder. (AM-138)*, pages 110–124. Princeton Univ. Press, 1995.
- [CW04] J. Crisp and B. Wiest. Embeddings of graph braid and surface groups in right-angled Artin groups and braid groups. *Algebr. Geom. Topol.*, 4:439–472, 2004.
- [Gro87] Misha Gromov. Hyperbolic groups. In S. M. Gersten, editor, *Essays in Group Theory, volume 8*. Springer-Verlag, New York, 1987.
- [HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geometric and Functional Analysis*, 17:1551–1620, 2008.
- [HW10] Tim Hsu and Daniel T. Wise. Cubulating graphs of free groups with cyclic edge groups. *American Journal of Mathematics*, 132(5):1153–1188, 2010.

- [KK13] Syang-hyun Kim and Thomas Koberda. Anti-trees and right-angled Artin subgroups of braid groups. *Preprint*, 2013. arXiv:1312.6465 [math.GR].
- [LS77] Roger Lyndon and Paul Schupp. *Combinatorial Group Theory*. Springer, 1977.
- [Mil68] John Milnor. A note on curvature and the fundamental group. *J. Differential Geom.*, 2:1–7, 1968.
- [Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proceedings of the London Mathematical Society*, 3(71(3)):585–617, 1995.
- [Sch27] Otto Schreier. Die untergruppen der freien gruppen. *Abhandlungen aus dem Mathematischen Seminar der Universitt Hamburg*, 5(1):161–183, 1927.
- [Scr10] Travis Scrimshaw. Embeddings of right-angled Artin groups. *Preprint*, 2010. arXiv:1010.6286 [math.GR].